

An Experimental Performance Comparison for Indexing Mobile Objects on the Plane

S. Sioutas
Department of Informatics
Ionian University
49100 Corfu, Greece
sioutas@ionio.gr

K. Tsihlias
Department of Informatics
Aristotle University of
Thessaloniki
54124 Thessaloniki, Greece
tsihlias@csd.auth.gr

G. Papaloukopoulos
Department of Computer
Engineering and Informatics
University of Patras
26500 Patras, Greece
papaluk@ceid.upatras.gr

Y. Manolopoulos
Department of Informatics
Aristotle University of
Thessaloniki
54124 Thessaloniki, Greece
manolopo@csd.auth.gr

ABSTRACT

We present a time-efficient approach to index objects moving on the plane to efficiently answer range queries about their future positions. Each object is moving with non small velocity u , meaning that the velocity value distribution is skewed (Zipf) towards u_{min} in some range $[u_{min}, u_{max}]$, where u_{min} is a positive lower threshold. Our algorithm enhances a previously described solution [18] by accommodating the ISB-tree access method as presented in [6]. Experimental evaluation shows the improved performance, scalability and efficiency of the new algorithm.

Categories and Subject Descriptors

H.2 [Database Management]: [Emergent Systems]

General Terms

Algorithms, Data Structures and Indexing, Experimentation

Keywords

Spatio-Temporal and Multimedia Databases, Data and Knowledge Management Systems

1. INTRODUCTION

This paper focuses on the problem of indexing mobile objects in two dimensions and efficiently answering range queries over the objects' future locations. This problem is motivated by a set of real-life applications such as intelligent transportation systems, cellular communications, and meteorology monitoring. The basic approach uses discrete

movements, where the problem of dealing with a set of moving objects can be considered as equivalent to a sequence of database snapshots of the object positions/extents taken at time instants $t_1 < t_2 < \dots$, with each time instant denoting the moment where a change took place. From this point of view, the indexing problems in such environments can be dealt with by suitably extending indexing techniques from the area of spatio-temporal databases [4, 17]. In [10] it is exposed how these indexing techniques can be generalized to handle efficiently queries in a discrete spatiotemporal environment. The common thrust behind these indexing structures lies in the idea of abstracting each object's position as a continuous function of time, $f(t)$, and updating the database whenever the function parameters change. Accordingly an object is modeled as a pair consisting of its extent at a reference time (design parameter) and of its motion vector. One categorization of the aforementioned structures is according to the family of the underlying access method used. In particular, there are approaches based either on R-trees or on Quadtrees as explained in [13, 14]. On the other hand, these structures can be also partitioned into those that: (a) are based on geometric duality and represent the stored objects in the dual space [1, 7, 12], and (b) leave the original representation intact by indexing data in their native dimensional space [3, 11, 15, 16, 19]. The *geometric duality transformation* is a tool extensively used in the Computational Geometry literature, which maps hyper-planes in R^d to points and vice-versa. In this paper we present and experimentally evaluate techniques using the duality transform as in [7, 11] to efficiently index future locations of moving points on the plane. In the next section we present a brief overview of the most basic practical methods. In Section 3 we give a formal description of the problem. In Section 4 we introduce the duality transform methods, in section 5 we briefly present our main contribution whereas in section 6 we present the ISBs access method that compares favorably with the solutions of [7, 11], the TPR* index [19], the STRIPES index [12] and the LBTs index [18] as well. In simple words, the new solution is the most efficient in terms of update I/O performance. Moreover, with respect to the query I/O performance, solution of ISBs is 4 or 5 faster than LBTs method

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MEDES 2009 October 27-30, 2009, Lyon, France

Copyright 2008 ACM 978-1-60558-829-2/08/0003 ...\$10.00.

and outperforms STRIPES (state of the art as of now) in many settings. Section 7 presents a thorough experimental evaluation, whereas Section 8 concludes the paper.

2. A BRIEF OVERVIEW OF THE RELEVANT METHODS

The TPR tree [15] in essence is an R^* -tree generalization to store and access linearly moving objects. The leaves of the structure store pairs with the position of the moving point and the moving point id, whereas internal nodes store pointers to subtrees with associated rectangles that minimally bound all moving points or other rectangles in the subtree. The difference with respect to the classical R^* -tree lies in the fact that the bounding rectangles are time parameterized (their coordinates are functions of time). It is considered that a time parameterized rectangle bounds all enclosed points or rectangles at all times not earlier than current time. Search and update algorithms in the TPR tree are straightforward generalizations of the respective algorithms in the R^* -tree; moreover, the various kinds of spatiotemporal queries can be handled uniformly in 1-, 2-, and 3-dimensional spaces.

The TPR-tree served as the base structure for further developments in the area [16]. *TPR*-tree*, an extension of the TPR-tree, improves the latter in update operations [19]. The main improvement lies in the fact that local optimization criteria (at each tree node) may degrade seriously the performance of the structure and more particularly in the use of update rules that are based on global optimization criteria. Thus, the authors of [19] proposed a novel probabilistic cost model to validate the performance of a spatiotemporal index and analyze with this model the optimal performance for any data-partition index.

The STRIPES index [12] is based on the application of the duality transformation and employs disjoint regular space partitions (disk based quadtrees [4]). Through the use of a series of implementations, the authors claim that STRIPES outperforms TPR^* -trees for both update and query operations.

Finally, the LBTs index [18] has the most efficient update performance in all cases. Regarding the query performance, LBTs method prevails as long as the query rectangle length remains in realistic levels (by far superiority in comparison to opponent methods). If the query rectangle length becomes huge in relation to the whole terrain, then STRIPES is the best solution, however, only to a small margin in comparison to LBTs method.

3. DEFINITIONS AND PROBLEM DESCRIPTION

We consider a database that records the position of moving objects in two dimensions on a finite terrain. We assume that objects move with a constant velocity vector starting from a specific location at a specific time instant. Thus, we can calculate the future object position, provided that its motion characteristics remain the same. Velocities are bounded by $[u_{min}, u_{max}]$. Objects update their motion information, when their speed or direction changes. The system is dynamic, i.e. objects may be deleted or new objects may be inserted.

Let $P_z(t_0) = [x_0, y_0]$ be the initial position at time t_0 of object z . If object z starts moving at time $t > t_0$, its position

will be $P_z(t) = [x(t), y(t)] = [x_0 + u_x(t - t_0), y_0 + u_y(t - t_0)]$, where $U = (u_x, u_y)$ is its velocity vector.

We would like to answer queries of the form: “Report the objects located inside the rectangle $[x_{1_q}, x_{2_q}] \times [y_{1_q}, y_{2_q}]$ at the time instants between t_{1_q} and t_{2_q} (where $t_{now} \leq t_{1_q} \leq t_{2_q}$), given the current motion information of all objects.”

4. INDEXING MOBILE OBJECTS USING DUALITY TRANSFORMATIONS

In general, the duality transform maps a hyper-plane h from R^d to a point in R^d and vice-versa. One duality transform for mapping the line with equation $y(t) = ut + a$ to a point in R^2 is by using the dual plane, where one axis represents the slope u of an object’s trajectory (i.e. velocity), whereas the other axis represents its intercept a . Thus we get the dual point (u, a) (this is the so-called *Hough-X transform* [7, 11]). By rewriting the equation $y = ut + a$ as $t = \frac{1}{u}y - \frac{a}{u}$, we arrive to a different dual representation (the so called *Hough-Y transform* in [7, 11]). The point in the dual plane has coordinates (b, w) , where $b = -\frac{a}{u}$ and $w = \frac{1}{u}$.

In [7, 11], motions with small velocities in the Hough-Y approach are mapped into dual points (b, w) having large w coordinates ($w = 1/u$). Thus, since few objects can have small velocities, by storing the Hough-Y dual points in an index such as an R^* -tree, Maximum Bounded Rectangles (MBRs) with large extents are introduced, and the index performance is severely affected. On the other hand, by using a Hough-X for the small velocities’ partition, this effect is eliminated, since the Hough-X dual transform maps an object’s motion to the (u, a) dual point. The query area in Hough-X plane is enlarged by the area E , which is easily computed as $E_{Hough-X} = (E1_{Hough-X} + E2_{Hough-X})$. By $Q_{Hough-X}$ we denote the actual area of the simplex query. Similarly, on the dual Hough-Y plane, $Q_{Hough-Y}$ denotes the actual area of the query, and $E_{Hough-Y}$ denotes the enlargement. According to these observations the solution in [7, 11] proposes the choice of that transformation which minimizes the criterion: $c = \frac{E_{Hough-X}}{Q_{Hough-X}} + \frac{E_{Hough-Y}}{Q_{Hough-Y}}$.

In order to build the index, we first decompose the 2-d motion into two 1-d motions on the (t, x) and (t, y) planes and then we build the corresponding index for each projection. Then we have to partition the objects according to their velocity: Objects with small velocity are stored using the Hough-X dual transform, whereas the rest are stored using the Hough-Y dual transform. Motion information about the other projection is also included.

To answer the exact 2-d query we decompose the query into two 1-d queries, for the (t, x) and (t, y) projection. Then, for each projection, we get the dual-simplex query and calculate the criterion c and choose the one (say p) that minimizes it. We search in projection p the Hough-X or Hough-Y partition and finally we perform a refinement or filtering step “on the fly”, by using the whole motion information. Thus, the result set contains only the objects satisfying the query.

5. OUR CONTRIBUTION

We consider the case, where the objects are moving with non small velocities u , meaning that the velocity value distribution is skewed (Zipf) towards u_{min} in some range $[u_{min}, u_{max}]$ and as a consequence the $Q_{Hough-Y}$ transformation is used (denote that u_{min} is a positive lower threshold). In [7, 11]

and [18], $Q_{Hough-Y}$ is computed by querying a B^+ -tree and LBTs (Lazy B-trees) respectively, each of which indexes the b parameters. Our construction is based on the use of the ISB-tree [6] instead of the B^+ -tree or Lazy B-trees, achieving optimal update performance and near-optimal query performance. Next we describe the main characteristics of the ISB-tree.

6. THE ISB-TREE

In the following, we give the required technical details of the external memory Interpolation Search B-Tree (ISB-tree) presented in [6]. First, we give some basic definitions about regular and smooth input distributions.




Figure 1: The ISB-tree Index

According to Willard [20], a probability density μ is regular if there are constants b_1, b_2, b_3, b_4 such that $\mu(x) = 0$ for $x < b_1$ or $x > b_2$, and $\mu(x) \geq b_3 > 0$ and $|\mu'(x)| \leq b_4$ for $b_1 \leq x \leq b_2$. This has been further pursued by Mehlhorn and Tsakalidis [9], who introduced the *smooth* input distributions, a notion that was further generalized and refined in [2]. Given two functions f_1 and f_2 , a density function $\mu = \mu[a, b](x)$ is (f_1, f_2) -smooth [2] if there exists a constant β , such that for all c_1, c_2, c_3 , $a \leq c_1 < c_2 < c_3 \leq b$, and all integers n , it holds that

$$\int_{c_2 - \frac{c_3 - c_1}{f_1(n)}}^{c_2} \mu[c_1, c_3](x) dx \leq \frac{\beta \cdot f_2(n)}{n}$$

where $\mu[c_1, c_3](x) = 0$ for $x < c_1$ or $x > c_3$, and $\mu[c_1, c_3](x) = \mu(x)/p$ for $c_1 \leq x \leq c_3$ where $p = \int_{c_1}^{c_3} \mu(x) dx$.

Intuitively, function f_1 partitions an arbitrary subinterval $[c_1, c_3] \subseteq [a, b]$ into f_1 equal parts, each of length $\frac{c_3 - c_1}{f_1} = O(\frac{1}{f_1})$; that is, f_1 measures how fine is the partitioning of an arbitrary subinterval. Function f_2 guarantees that no part, of the f_1 possible, gets more probability mass than $\frac{\beta \cdot f_2}{n}$; that is, f_2 measures the sparseness of any subinterval $[c_2 - \frac{c_3 - c_1}{f_1}, c_2] \subseteq [c_1, c_3]$. The class of (f_1, f_2) -smooth distributions (for appropriate choices of f_1 and f_2) is a superset

of both regular and uniform classes of distributions, as well as of several non-uniform classes [2, 5]. Actually, *any* probability distribution is $(f_1, \Theta(n))$ -smooth, for a suitable choice of β .

The ISB-tree is a two-level data structure (see Figure 1). The upper level is a non-straightforward externalization of the Static Interpolation Search Tree (SIST) presented in [5]. A static interpolation search tree (SIST) can be fully characterized by three non-decreasing functions $H(n)$, $R(n)$ and $I(n)$, where $H(n)$ denotes the tree height, $R(n)$ denotes the root fan-out and $I(n)$ denotes how fine is the partition of the set of elements. If in the basic searching routine presented in Algorithm 1, we set $next \leftarrow \lfloor \frac{x - S[left]}{S[right] - S[left]}(right - left) \rfloor + left$ we refer to an interpolation searching routine for which time improvements can be obtained if certain classes of *smooth* input distributions are considered. A natural extension is to adapt interpolation search into dynamic data structures, that is, the method above (the so-called **Static Interpolation Search Tree**).

Algorithm 1 Search(x, S)

```

1: left ← 1
2: right ← n
3: next ← k ∈ [left, right]
4: while x <> S[next] and left < right do
5:   if x ≤ S[next] then
6:     right ← next - 1
7:   else
8:     left ← next + 1
9:     next ← k ∈ [left, right]
10:  end if
11: end while
12: if x = S[next] then
13:   print('Success')
14: else
15:   print('Fail')
16: end if

```

The lower level of the ISB-tree is a forest of buckets, each one implemented by a new variant of the classical B-tree, the *Lazy B-tree*, introduced in [6] and used in [18]. Each bucket contains a subset of the stored elements and is represented by a unique representative. The representatives of the buckets as well as some additional elements are stored in the upper level structure. The following theorem provides the complexities of the Lazy B-tree:

Theorem 1. The Lazy B-Tree supports the search operation in $O(\log_B N)$ worst-case block transfers and update operations in $O(1)$ worst-case block transfers, provided that the update position is given.

Proof. See [6, 18].

The upper level data structure is an external version T of the static interpolation search tree (SIST) [5], with parameters $R(s) = s^\delta$, $I(s) = s/(\log \log s)^{1+\epsilon}$, where $\epsilon > 0$, $\delta = 1 - \frac{1}{B}$, and s is the number of stored elements in the tree. The specific choice of δ guarantees the desirable $O(\log_B \log s)$ height of the upper level structure. For each node that stores more than $B^{1+\frac{1}{B-1}}$ elements in its subtree, we represent its REP and ID arrays as static external sorted arrays; otherwise, we store all the elements in a constant number of disk blocks. In particular, let v be a node and n_v be the number of stored elements in its subtree,

with $n_v \geq B^{1+\frac{1}{B-1}}$. Node v is associated with two external arrays $EREP_v$ and EID_v that represent the REP_v and ID_v arrays of the original SIST structure. The EID_v array uses $O\left(\frac{I(n_v)}{B}\right)$ contiguous blocks, the $EREP_v$ array uses $O\left(\frac{R(n_v)}{B}\right)$ contiguous blocks, while an arbitrary element of the arrays can be accessed with $O(1)$ block transfers, given its index. Moreover, the choice of the parameter $B^{1+\frac{1}{B-1}}$ guarantees that each of the $EREP_v$ and EID_v arrays contains at least B elements, and hence we do not waste space (in terms of underfull blocks) in the external memory representation.

To insert/delete an element, given the position (block) of the update, we simply have to insert/delete the element to/from the Lazy B-tree storing the elements of the corresponding bucket. Note that the external SIST is not affected by these updates. Each time the number of updates exceeds cn_0 , where c is a predefined constant, the whole data structure is reconstructed.

The search procedure for locating a query element x can be decomposed into two phases: (i) the traversal of internal nodes of the external SIST to locate a bucket \mathcal{B}_i , and (ii) the search for x in the Lazy B-tree storing the elements of \mathcal{B}_i . For more technical details see [6]. Algorithms 2-5 provide the description of ISB-tree's basic operations in pseudocode.

Lemma 2: The traversal of internal nodes of the external SIST requires $O(\log_B \log n)$ expected block transfers with high probability.

Proof: See [6].

The insertions and deletions of elements into the ISB-tree were simulated by a combinatorial game of balls and bins described in [5] for an internal finger-search data structure. In particular, balls correspond to elements and bins to buckets. Insertions of elements into the ISB-tree were simulated by the insertion of balls into bins according to an unknown smooth probability density μ . Similarly, the deletion of an element from the ISB-tree was simulated by the deletion of an element from a bin uniformly at random. For this process the following has been proven in [5].

Algorithm 2 Sist_Search($x, SIST$)

```

1:  $v \leftarrow root(SIST)$ 
2: while ( $v \llcorner leaf$ ) do
3:    $i = \lfloor \frac{x-l_v}{u_v-l_v} R(n_v) \rfloor$  {Let  $v$  be a node on the search path for  $x, n_v$  the number of leaves in its subtree,  $l_v$  and  $u_v$  the minimum and maximum element respectively, stored in  $T_v$ }
4:   Retrieve the  $\lceil \frac{i}{B} \rceil$ -th block of the  $EID_v$  array
5:    $j = EID_v[i]$ 
6:    $l = \lceil \frac{j}{B} \rceil$ 
7:   repeat
8:      $l \leftarrow l + 1$ 
9:   until  $EREP_v[l] \leq x < EREP_v[l + 1]$ 
10: end while
11: follow the pointer from leaf  $v$  {Let  $Bin_i$  the corresponding bin which is organized as a Lazy B-tree}

```

Theorem 2: Consider the combinatorial game of balls and bins described in [5]. Then, the expected number of balls in a bin is $O(\log n)$ with high probability.

The following lemma establishes the searching bound within a bucket of the ISB-tree.

Algorithm 3 ISB_Search(x, ISB -tree)

```

1: Sist_Search( $x, SIST$ ) {Let  $Bin_i$  the corresponding Bin of the static interpolation search tree SIST}
2: LazyTree_Search( $x, Bin_i$ ) {search in the lazy B-tree was implemented in [18]}

```

Lemma 3: Searching for an element in a bucket of the ISB-tree takes $O(\log_B \log n)$ expected block transfers with high probability.

Proof: This is an immediate result from Theorem 1 and the size of each bucket, which is determined by Theorem 2.

The main theorem presented in [6] follows and holds for the very broad class of $(n/(\log \log n)^{1+\epsilon}, n^{1-\delta})$ -smooth densities, where $\delta = 1 - \frac{1}{B}$ and includes the uniform, regular, bounded as well as several non-uniform distributions.

Algorithm 4 ISB_Insert(x, ISB -tree)

```

1: Sist_Search( $x, SIST$ ) {Let  $Bin_i$  the corresponding Bin}
2: LazyTree_Insert( $x, Bin_i$ ) { $Bin_i$  is a lazy B-tree and its insert operation was implemented in [18]}
3:  $numberofupdates \leftarrow numberofupdates + 1$ 
4: if  $numberofupdates = cn_0$  then
5:   Rebuild( $SIST$ ) { $n_0$  is the total number of elements stored in the initial ISB-tree}
6: end if

```

Algorithm 5 ISB_Delete(x, ISB -tree)

```

1: Sist_Search( $x, SIST$ ) {Let  $Bin_i$  the corresponding Bin}
2: LazyTree_Delete( $x, Bin_i$ ) { $Bin_i$  is a lazy B-tree and its delete operation was implemented in [18]}
3:  $numberofupdates \leftarrow numberofupdates + 1$ 
4: if  $numberofupdates = cn_0$  then
5:   Rebuild( $SIST$ ) { $n_0$  is the total number of elements stored in the ISB-tree at the initial state}
6: end if

```

Theorem 3: Suppose that the upper level of the ISB-tree is an external static interpolation search tree with parameters $R(s_0) = s_0^\delta$, $I(s_0) = s_0/(\log \log s_0)^{1+\epsilon}$, where $\epsilon > 0$, $\delta = 1 - \frac{1}{B}$, $s_0 = n_0$, n_0 is the number of elements in the latest reconstruction, and that the lower level is implemented as a forest of Lazy B-trees. Then, the ISB-tree supports search operations in $O(\log_B \log n)$ expected block transfers with high probability, where n denotes the current number of elements, and update operations in $O(1)$ worst-case block transfers, if the update position is given. The worst-case update bound is $O(\log_B n)$ block transfers, and the structure occupies $O(n/B)$ blocks.

Proof¹: From Lemmas 2 and 3, the searching operation takes $O(\log_B \log n)$ expected number of block transfers with high probability. Considering the update bound, between reconstructions the block transfers for an update are clearly $O(1)$, since we only have to update the appropriate Lazy B-tree which can be done in $O(1)$ block transfers (see Theorem 1). The reconstructions can be easily handled by using the technique of global rebuilding [8]. With this technique the linear work spent during a global reconstruction of the upper level structure may be spread out on the updates in such a

¹We quote a brief description of the proof presented in [6]

way that a rebuilding cost of $O(1)$ block transfers is spent at each update.

7. EXPERIMENTAL EVALUATION

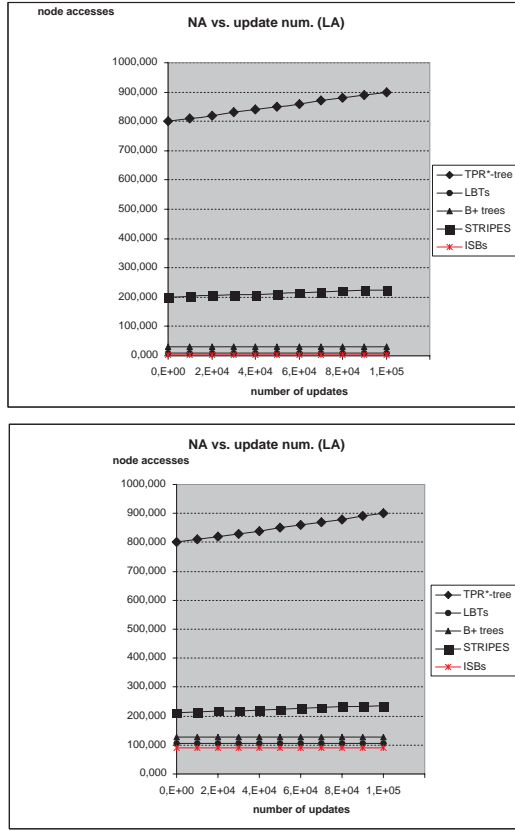


Figure 2: $q_vlen = 5, q_Tlen = 50, q_Rlen = 100$ (top), $q_Rlen = 1000$ (bottom)

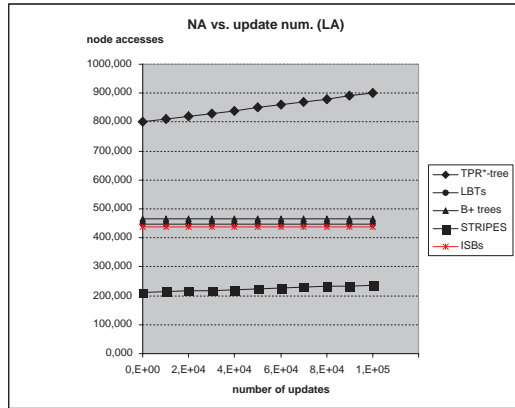


Figure 3: $q_Rlen = 2000, q_vlen = 5, q_Tlen = 50$

This section compares the query/update performance of our solution with STRIPES as well as with those ones that use B⁺-trees, Lazy B-trees (LBTs) and TPR*-tree, as well. We deploy spatio-temporal data that contain insertions at a

single timestamp 0. In particular, objects' MBRs are taken from the LA spatial dataset². We want to simulate a situation where all objects move in a space with dimensions 100x100 kilometers. For this purpose each axis of the space is normalized to [0,100000]. For the TPR*-tree, each object is associated with a VBR (Velocity Bounded Rectangle) such that (a) the object does not change spatial extents during its movement, and (b) the velocity value distribution is skewed (Zipf) towards 30 in the range [30,50], and (c) the velocity can be either positive or negative with equal probability.

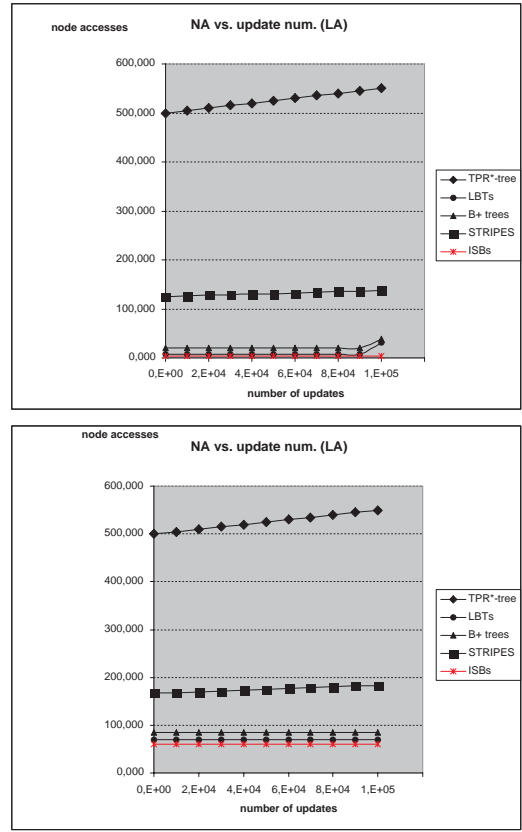


Figure 4: $q_vlen = 10, q_Tlen = 50, q_Rlen = 400$ (top), $q_Rlen = 1000$ (bottom)

We will use a page size of 1 Kbyte so that the number of index nodes simulates realistic situations. Also, for all experiments, the key length is 8 bytes, whereas the pointer length is 4 bytes. Thus, the maximum number of entries ($\langle x \rangle$ or $\langle y \rangle$, respectively) in both Lazy B-trees and B⁺-trees is $1024/(8+4)=85$. In the same way, the maximum number of entries (2-d rectangles or $\langle x_1, y_1, x_2, y_2 \rangle$ tuples) in TPR*-tree is $1024/(4*8+4)=27$. On the other hand, the STRIPES index maps predicted positions to points in a dual transformed space and indexes this space using a disjoint regular partitioning of space. Each of the two dual partitioning, are equally partitioned into four quads. This partitioning results in a total of $4^2 = 16$ partitions, which we call *grids*. Thus, the fan-out of each internal node is 16. The ISB-tree has an exponentially decreased fan-out and 2 levels

²Downloaded 128.971 MBRs from <http://www.census.gov/geo/www/tiger/>

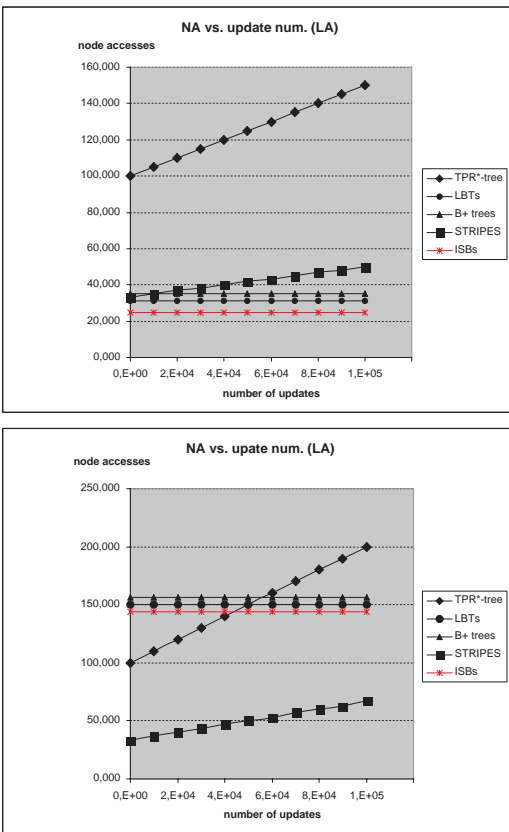


Figure 5: $q_vlen = 5$, $q_Tlen = 1$, $q_Rlen = 400$ (top), $q_Rlen = 1000$ (bottom)

at most (source code of ISB-tree access method is available at <http://www.ionio.gr/~sioutas/New-Software.htm>).

For each dataset, all indexes except for STRIPES and ISBs have similar sizes. Specifically, for LA, each tree has 4 levels and around 6700 leaves with the exception of: (a) the STRIPES index which has a maximum height of seven and consumes about 2.4 times larger disk space, and (b) the ISB index which has a maximum height of 2. Each query q has three parameters: q_Rlen , q_vlen , and q_Tlen , such that: (a) its MBR q_R is a square, with length q_Rlen , uniformly generated in the data space, (b) its VBR is $q_v = [-q_vlen/2, q_vlen/2, -q_vlen/2, q_vlen/2]$, and (c) its query interval is $q_T = [0, q_Tlen]$. The query cost is measured as the average number of node accesses in executing a workload of 200 queries with the same parameters. Implementations were carried out in C++ including particular libraries from SECONDARY LEDA v4.1.

7.1 Query cost comparison

We measure the performance of our technique described previously (in particular one ISB-tree for each projection, plus the query processing between the two answers), in comparison to that of the LBTs method [18], the traditional technique presented in [7, 11], the TPR*-tree [19] and the STRIPES method [12], using the same query workload, after every 10000 updates. Figures 2-6 show the query cost (for datasets generated from LA as described above) as a function of the number of updates, using workloads with

different parameters. In these figures our solution is almost 4-5 times more efficient (in terms of the number of I/Os) than the solution using LBTs and B^+ -trees. This fact is an immediate result of the sublogarithmic I/O searching complexity of ISB-tree in comparison to the logarithmic I/O searching complexities of both structures B^+ -tree and Lazy B-tree. In particular, we have to index the appropriate b parameters in each projection and then to combine the two answers by detecting and filtering all pair permutations. As a consequence, the ISBs method is significantly faster than LBTs and traditional B^+ -trees methods.

Figure 2 depicts the efficiency of our solution in comparison to that of the TPR*-tree and STRIPES. In the top figure, where the length of the query rectangle is 100 and as a consequence the query's surface is equal to $10000m^2$ or 1 hectare (the surface of the whole spatial terrain is 10^6 hectares) the ISBs method is consistently about 53 times faster than the STRIPES method, 212 times faster than the TPR*-tree, 7.5 times faster than the B^+ -trees method and 2 times faster than the LBTs method. The superiority of our solution decreases as the query rectangle length grows from 100 to 1000. Thus, in the bottom figure, where the spatial query's surface is equal to 100 hectares, again our method is faster about 2.2 times with respect to the STRIPES method, 8.3 times wrt the TPR*-tree, 1.25 times wrt the B^+ -trees methods and 1.05 times wrt the LBTs method.

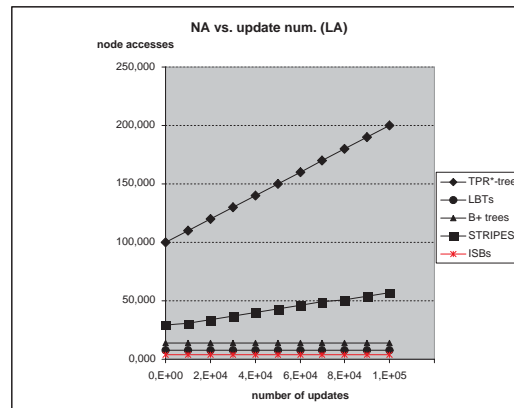


Figure 6: $q_Rlen = 400$, $q_vlen = 5$, $q_Tlen = 100$

In real GIS applications, for a vast spatial terrain of 10^6 hectares, e.g. the road network of a big town where each road square covers no more than 1 hectare (or $10.000m^2$) the most frequent queries consider spatial query's surface no more than 100 road squares (or 100 hectares) and future time interval no larger than 100 seconds. However, to test the methods' performance in extreme cases we conducted the following experiment. When the query rectangle length or equivalently the query surface becomes extremely large (e.g. 2000, or equivalently 400 hectares), then the STRIPES index shows better performance as depicted in Figure 3. In particular, our method is still 1.9 times faster than the TPR*-tree, however, the STRIPES method is twice faster than our one. The apparent explanation is that as the surface of the query rectangle grows, the answer size in each projection grows as well, thus the performance of the ISBs method that combines and filters the two answers becomes less attractive. However, we do not consider such extreme

case as *realistic* scenarios. Figure 4 depicts the performance of all methods for a growing velocity vector. In particular, in the top figure the ISBs method consistently prevails about 33 times in comparison to the STRIPES method, 137 times in comparison to the TPR*-tree, 5 times in comparison to the B⁺-trees and 2 times in comparison to the LBTs method. The superiority of our solution becomes less strong as the query rectangle length grows from 400 (16 hectares of query surface) to 1000 (100 hectares of query surface). However, notably even in the latter case (see bottom of the figure), our method is about 2.7 times faster with respect to the STRIPES method, 8.3 times wrt the TPR*-tree, 1.3 times wrt the B⁺-trees and 1.06 times wrt the LBTs method. Obviously, the velocity factor is very important for TPR-like solutions, but not for the other methods, for LBTs and ISBs in particular, which depend exclusively on the query surface. Figure 5 depicts the performance of all methods when the time interval length approaches the 1 value. However, notably even in this case (see top of the figure), the ISBs method is about 1.6 times faster with respect to the STRIPES method, 4.6 times faster wrt the TPR*-tree, 1.3 times faster wrt the B⁺-trees and 1.2 times faster wrt the LBTs method. As query rectangle length grows from 400 to 1000, the ISBs method advantage decreases; from the bottom figure, we remark that STRIPES is about 3 times faster, whereas our method is 1.03 times faster than the TPR*-tree, 1.07 times faster than the B⁺-trees and 1.03 times faster than the LBTs method. Figure 6 depicts the efficiency of our solution in comparison to that of TPR*-trees and STRIPES when the time interval length reaches the value of 100. In particular, the ISBs method is consistently about 10 times faster than STRIPES, 37 times faster than TPR*-tree, 3.5 times faster than the B⁺-trees and 2 times faster than the LBTs method. As required in practice, the query surface remains in *realistic* levels (16 hectares).




Figure 7: Update Cost Comparison

7.2 Scalability and update cost comparison

Figure 7 compares the average cost (amortized over insertions and deletions) as a function of the number of updates. ISBs and LBTs methods have optimal update performance and consistently outperform the TPR*-tree by a wide margin as well as the STRIPES index by a narrow margin. In particular, ISBs and LBTs methods require a constant number of 4 and 6 block transfers respectively (3 and 2 block

transfers respectively for each projection, for details see [6]) and this update performance is independent of the dataset size. On the other hand, the other 3 solutions do not have constant update performance; instead their performance depends on the dataset size even if as in the experiment of Figure 7 B⁺-trees and STRIPES reach the optimal performance of ISBs and LBTs methods requiring 8 and 7 block transfers respectively (TPR*-tree requires 35 block transfers in average). The experiments above show that ISBs method achieves a near optimal performance for the most cases. This stems from the fact that the MBRs’ projections from the LA spatial datasets follow an almost uniform (the most popular density of smooth family) distribution, due to the almost uniform decomposition of spatial maps. In particular, LA dataset constitutes a dense spatial map and hence the derived one-dimensional data produce densely populated elements. Thus, the interpolation technique of ISBs method works very well and its expected excellent behavior follows with high probability.

8. CONCLUSIONS

We have used a new access method for indexing moving objects on the plane to efficiently answer range queries about their future location. Its update performance is the most efficient in all cases with no exception. Regarding the query performance, the superiority of our structure has been shown as long as the query rectangle length remains in realistic levels, thus by far outperforming the opponent methods. If the query rectangle length becomes extremely huge in relation to the whole terrain (which apparently is a non-practical instance), then the STRIPES method is the best solution, however, only to a small margin in comparison to our method. We anticipate that for synthetic gigantic datasets the ISBs method will be superior in any case.

9. REFERENCES

- [1] Agarwal P.K., Arge L. and Erickson J.: “Indexing Moving Points”, *Proceedings 19th ACM Symposium on Principles of Database Systems (PODS)*, pp.175-186, Dallas, TX, 2000.
- [2] Andersson A. and Mattson C.: “Dynamic Interpolation Search in $o(\log \log n)$ Time”, *Proceedings 20th International Colloquium on Automata, Languages and Programming (ICALP)*, pp.15-27, Lund, Sweden, 1993.
- [3] Beckmann N., Kriegel H.P., Schneider R. and Seeger B.: “The R*-tree: an Efficient and Robust Access Method for Points and Rectangles”, *Proceedings ACM International Conference on Management of Data (SIGMOD)*, pp.322-331, Atlantic City, NJ, 1990.
- [4] Gaede V. and Gunther O. “Multidimensional Access Methods”, *ACM Computing Surveys*, Vol.30, No.2, pp.170-231, 1998.
- [5] Kaporis A., Makris C., Sioutas S., Tsakalidis A., Tsihlias K. and Zaroliagis C.: “Improved Bounds for Finger Search on a RAM”, *Proceedings 12th Annual European Symposium on Algorithms (ESA)* pp.325-336, Bergen, Norway, 2003. Full version as Technical Report TR-2003/07/01, Computer Technology Institute, Patras, 2003.
- [6] Kaporis A., Makris C., Sioutas S., Tsakalidis A., Tsihlias K. and Zaroliagis K.: “ISB-Tree: A New Indexing Scheme with Efficient Expected Behaviour”,

- Proceedings 13th International Symposium on Algorithms and Computation (ISAAC)*, pp.318-327, Sanya, Hainan, China, 2005.
- [7] Kollios G., Gunopulos D. and Tsotras V.: "On Indexing Mobile Objects", *Proceedings 18th ACM Symposium on Principles of Database Systems (PODS)*, pp.261-272, Philadelphia, PA, 1999.
- [8] Levkopoulos S. and Overmars M.H.: "Balanced Search Tree with $O(1)$ Worst-case Update Time", *Acta Informatica*, Vol.26, No.3, pp.269-277, 1988.
- [9] Mehlhorn K. and Tsakalidis A.: "Dynamic Interpolation Search", *Journal of the ACM*, Vol.40, No.3, pp.621-634, 1993.
- [10] Manolopoulos Y., Theodoridis Y. and Tsotras V.: "*Advanced Database Indexing*", Kluwer Academic Publishers, 2000.
- [11] Papadopoulos D., Kollios G., Gunopulos D. and Tsotras V.J.: "Indexing Mobile Objects on the Plane", *Proceedings 13th International Workshop on Database and Expert Systems Applications (DEXA)*, pp.693-697, Aix-en-Provence, France, 2002.
- [12] Patel J., Chen Y. and Chakka V.: "STRIPES: an Efficient Index for Predicted Trajectories", *Proceedings ACM International Conference on Management of Data (SIGMOD)*, pp.637-646, Paris, France, 2004.
- [13] Raptopoulou K., Vassilakopoulos M. and Manolopoulos Y.: "Towards Quadtree-based Moving Objects Databases", *Proceedings 8th East-European Conference on Advanced Databases and Information Systems (ADBIS)*, pp.230-245, Budapest, Hungary, 2004.
- [14] Raptopoulou K., Vassilakopoulos M. and Manolopoulos Y.: "On Past-time Indexing of Moving Objects", *Journal of Systems and Software*, Vol.79, No.8, pp.1079-1091, 2006.
- [15] Saltenis S., Jensen C., Leutenegger S. and Lopez M.A.: "Indexing the Positions of Continuously Moving Objects", *Proceedings ACM International Conference on Management of Data (SIGMOD)*, pp.331-342, Dallas, TX, 2000.
- [16] Saltenis S. and Jensen C.S.: "Indexing of Moving Objects for Location-Based Services", *Proceedings 18th IEEE International Conference on Data Engineering (ICDE)*, pp.463-472, San Jose, CA, 2002.
- [17] Salzberg B. and Tsotras V.J.: "A Comparison of Access Methods for Time-Evolving Data", *ACM Computing Surveys*, Vol.31, No.2, pp.158-221, 1999.
- [18] Sioutas S., Makris C., Tsihlias K., Tsakalidis K. and Manolopoulos Y.: "Indexing Mobile Objects on the Plane Revisited", *Proceedings 11th East European Conference on Advances in Databases and Information Systems (ADBIS)*, pp.189-204, Varna, Bulgaria, 2007
- [19] Tao Y., Papadias D. and Sun J.: "The TPR*-tree: an Optimized Spatio-Temporal Access Method for Predictive Queries", *Proceedings 29th International Conference on Very Large Data Bases (VLDB)*, pp.790-801, Berlin, Germany, 2003.
- [20] Willard D.E.: "Searching Unindexed and Nonuniformly Generated Files in $\log \log N$ Time", *SIAM Journal of Computing*, Vol.14, No.4, pp.1013-1029, 1985.